



THE VALUE OF INFORMATION IN DISTRIBUTED DECISION NETWORKS

Munther Dahleh
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

03/04/2016
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA2
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE					<i>Form Approved OMB No. 0704-0188</i>	
<small>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</small>						
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.						
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE			3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)					8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT						
15. SUBJECT TERMS						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)	

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATES COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report, e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5d. PROJECT NUMBER. Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

The Value of Information in Distributed Decision Networks
Grant No. FA9550-11-1-0312

Munther A. Dahleh

EECS Department and
Laboratory for Information and Decision Systems, MIT

Research Accomplishments

Below, we describe the various accomplishment of our work under grant FA9550-11-1-0312. We start by describing a general formulation, and then we describe the various results attained.

1 Mathematical description of Distributed Decision Network under Information Constraints

We now define a mathematical framework for networks. Let $\mathcal{G} = (V, E)$ be an undirected random network (graph) drawn from a known distribution $p_{\mathcal{G}}$,¹ composed of a finite vertex set V and a link set $E = \{V \times V \text{ modulo } S\}$ where $S = \{\{(i, j), (j, i)\}\}$. Each vertex $i \in V$ corresponds to an agent, and each link $(j, i) \in E$ corresponds to a channel by which information flows from agent j to agent i in the network. We denote the neighborhood of i by $\mathcal{N}(i) = \{j \mid (i, j) \in E\}$.

There is a state (internal or external) W drawn from a distribution p_W that the agents may want to estimate, transmit, and act upon. Each agent i also possesses a state and some private observation about W . We denote the state at time t by $x_i(t)$, and we assume that the tuple of initial states $(x_i(0))$ is correlated with W and are drawn randomly from a joint distribution p_{WX_0} . Agent i 's private information/observation at time t is denoted by $Y_i(t)$ and has a joint distribution $p_{WY_i(t)}$ with W . Finally, agent i has some information about agent j 's state (either because it can observe it or agent j transmits it), which we denote by $m_{ji}(t) = m_{ji}(x_j(t))$.

The agents autonomously update their states according to the dynamics

$$x_i(t+1) = f_i(x_i(t), \{m_{ji}(t)\}_{j \in \mathcal{N}(i)}, Y_i(t)). \quad (1)$$

The performance of the system, denoted as $J(\{f_i\}, \{m_{ji}\}, \{p_{WY_i}\}, p_{WX_0})$, is governed by the information $\{p_{WY_i}\}$ and the dynamics/algorithms $\{f_i\}$ and $\{m_{ji}\}$, and so we can consider optimizing the system's performance over these parameters. To this end, we consider a class of information types $\{p_{WY_i}\} \in \mathcal{P}$ over which the information can take its distribution as well as a set of dynamics/algorithms $\{f_i\} \in \mathcal{F}$ and messages/information $\{m_{ji}\} \in \mathcal{M}$ that appropriately constrain the dynamics of the system.

2 Static Decision Networks under Communication Constraints

Below, we summarize known research and our main contributions in the static problem.

2.1 Hypothesis Testing under Capacity Constraints

One of the simplest decision systems one can consider is binary hypothesis testing, where a decision between two hypothesis is made using observations. This fundamental problem has application in, for example, target identification and multi-mode systems identification.

¹At times, we will be interested in analyzing the performance of the network with respect to broad properties determined by $p_{\mathcal{G}}$ while at other times we will be interested in an analysis with respect to properties of a specific graph G . In the latter case, we simply set $p_{\mathcal{G}}$ to the degenerate distribution $p_{\mathcal{G}}(\mathcal{G}) = 1_{\mathcal{G}=G}$.

In this problem, the vertex set V is a single agent gathering information Y about a random variable X through a channel of fixed capacity c . The state of the world W can assume two discrete states w_1 and w_2 with a known probabilities, and the distribution of X depends on W ; that is, $X \sim p_i$ if $W = w_i$ where p_i belongs to the n -th dimensional simplex. The information $(Y(t))_t$ gathered is a sequence of i.i.d samples which have been filtered through the channel p_{WY} and whose alphabet has a size k that can be chosen arbitrarily.

In this setting, we use the traditional Shannon channel capacity parametrization of imperfect information. Specifically, we use sets $\mathcal{P}(c)$ parametrized by a capacity c , according to $\mathcal{P}(c) = \{p_{XY} \mid \max_{p_Y} I(X; Y) \leq c\}$. We also use the following approximation valid for small capacities derived in [?]:

$$\mathcal{P}'(c, p_0) = \left\{ p_{XY} \mid \frac{1}{2} \|p_{Y|X=i} - p_0\|_{[p_0^{-1}]}^2 \leq c \quad i = 1, \dots, n \right\},$$

where the capacity constraint is replaced by n quadratic inequalities and the output probability distribution lies around p_0 .

In this framework, the agent's performance is naturally based on the probability of error; that is, declaring " $W = w_1$ " when the true state of the world is $W = w_2$ and vice versa. In our case, rather than measure performance by the minimum number of samples required to make a decision with a fixed probability of error (quickest detection), we focus our attention on the rate at which the probability of error decreases as the number of samples collected increases. In this setting, the relationship between information and detection rates (i.e., the *value of information*) is related to a function $C(p_1, p_2, p_{XY})$ called the *Chernoff information*. The Chernoff information relates the asymptotic behavior of the estimation to the probability of error, and for a large number of samples t , the probability of error is written in terms the Chernoff information as

$$P_e(t) = e^{-C(p_1, p_2, p_{XY}) t}.$$

Therefore, in our case, we can equivalently express the average performance of the system with a channel p_{XY} as $J(p_{XY}) = e^{-C(p_1, p_2, p_{XY})}$ so that the optimal performance of such a system over all possible channels is obtained by solving the optimization problem

$$J(c) = \min_{p_{XY} \in \mathcal{P}(c)} J(p_{XY}). \quad (2)$$

We can explicitly solve this problem in the small capacity regime by substituting $\mathcal{P}(c)$ with $\mathcal{P}'(c)$ and by using quadratic approximation for the Chernoff Information, yielding the following theorem [?].

Theorem 1 (Value of information in hypothesis testing) *For small capacities c ,*

$$J(c) = 1 - \frac{\|p_1 - p_2\|_1^2}{4} c.$$

The optimizing distribution p_{XY}^* to optimization (??) that yields Theorem ?? also allows us to get an idea of the *actionable information* in this framework. Let $X_1 = \{x \mid p_1(x) > p_2(x)\}$ and $X_2 = \{x \mid p_2(x) > p_1(x)\}$. The optimizing channel applies opposing weights to symbols

in X_1 and X_2 , and it applies an arbitrary weights to symbols not in $X_1 \cup X_2$, the so-called *inactionable information* set. This is an important point for this application – a channel optimized for estimation of X is not necessarily the channel optimized for hypothesis estimation.

We also have the following relationship between the size of the output alphabet and the channel capacity.

Theorem 2 (Information content and latency in hypothesis testing) *Let k be the size of the output alphabet (i.e., $Y \in \{y_i\}_{1 \leq i \leq k}$) and let p_{XY}^k be the optimizing channel for that output alphabet size. For $k \geq k_{\min} = \lceil e^c \rceil$, $J(p_{XY}^k) = J(p_{XY}^{k_{\min}})$.*

In other words, for a fixed capacity, optimizing over a set of channels with output alphabet greater than $\lceil e^c \rceil$ does not improve performance. Another interesting aspect of the optimizing channels is that because they are tuned to hypothesis testing rather than estimation of X , decoding information through the channel is simple.

2.2 The Impact of the Network Topology on Distributed Hypothesis Testing

We now move to a type of distributed hypothesis testing where the agents seek to guess the state of the world W through individual trials. However, rather than study how the error rate diminishes as an agent accumulates unbounded information, we study how information flow through the networks impacts the error rate over an unbounded chain of agents. Ultimately, we are interesting in determining if the network eventually “guesses” the correct state of the world, in which case we say that the network “learned.” Interestingly, learning is not guaranteed in this setting. As we will see, the decision network’s topology can have unintended effects on how errors propagate through the network.

To illustrate the theoretical issues arising in this context, let us consider a learning problem over a social network of rational agents. The problem is whether agents would be able to extract valuable information about an unobservable parameter, simply by observing the behavior of their neighbors. In other words, we are interested in understanding which network structures would lead to *learning* the unknown parameter and which networks can generate a *herd behavior*, where in the Perfect Bayesian Nash Equilibrium, agents can only extract limited amount of information from their neighbors’ actions.

As before, let W denote an underlying state of the world, unknown to the agents, and suppose for simplicity that W takes two values W_L and $W_H > W_L$, and suppose that each individual receives an imperfectly informative signal about the value of W (also referred to as his *private belief*), denoted by Y , which is identical and independently distributed across individuals. It is common knowledge that the signal has a conditional distribution $p(Y|W)$ in states $W = W_H$ and $W = W_L$. We assume that the distributions $p(Y|W_H)$ and $p(Y|W_L)$ are absolutely continuous with respect to one another and have a common support $[\tilde{\sigma}, \bar{\sigma}] \subseteq [[0, 1]$. We say that the private beliefs are *bounded* if $0 < \tilde{\sigma} < \bar{\sigma} < 1$; and *unbounded* if $[\tilde{\sigma}, \bar{\sigma}] = [0, 1]$. For unbounded beliefs, signals in favor of state W_L are more likely to occur in state W_L than in state W_H , i.e., there is an underlying tendency for the truth to be revealed in the signals.

A realistic framework of learning in a multi-agent system must model structure of social networks with which individuals observe and communicate with each other.² Thus, we assume that individuals form a social network in which each agent can observe the actions of (a subset of) other individuals that have moved in the past. After observing their private signal (Y_i) and other available information about the actions of other individuals, each individual makes a decision. Beliefs are formed in a Bayesian manner based on the content of the private signal and the knowledge of what all precedents have done (i.e., actions of people who have moved before).

We are interested in understanding what agents can learn from one another in the long run, when they update their beliefs in a Bayesian fashion. To this end, we develop a theory of learning and dynamic belief formation when individuals observe not the entire action history, but rather the actions of a neighborhood of individuals. Notice that both the analysis and the equilibrium outcome are significantly different if individuals do not observe all past actions, but only a subset of these past actions, that may be for example randomly chosen from the entire set of past actions. One difficulty with this class of models is that to determine how beliefs will evolve, we need to characterize the perfect Bayesian Nash equilibrium, which involves rather complex inferences by individuals.

In recent work [?], we have developed a new framework for learning dynamics over a very general (deterministic or stochastic) social network of agents (see also related work [?] and [?]). In particular, we consider a countably infinite number of agents, each of which is making a decision x_n sequentially. We assume that the neighborhood of agent n , $B(n)$, is stochastically generated according to an arbitrary probability distribution p_G^n over the set of all subsets of $\{1, \dots, n-1\}$. The sequence $\{p_G^n\}$ is the *network topology* of the social network formed by the agents. The network topology is common knowledge, whereas the realized neighborhood $B(n)$ is the private information of agent n . Notice that in the case that $B(n)$ is a strict subset of $\{1, \dots, n-1\}$ for some $n \geq 2$, then the social beliefs do not form a martingale, and as a result, one cannot apply Doob's martingale convergence theorem in the analysis.

We provide a systematic characterization of the conditions under which there will be equilibrium information aggregation in social networks. We say that there is *information aggregation* or equivalently *asymptotic learning*, when, in the limit as the size of the social network becomes arbitrarily large, individual actions converge (in probability) to the action that yields the higher payoff. The key property of the network topology relevant to asymptotic learning turns out to be the *expanding observations property*.

To describe this concept, let us first introduce another notion: a finite group of agents is *excessively influential* if there exists an infinite number of agents who, with probability uniformly bounded away from 0, observe only the actions of a subset of this group. For example, a group is excessively influential if it is the source of *all* information (except individual signals) for an infinitely large component of the social network. If there exists an

²Although there is a large literature in economics on social learning (see [?], [?], [?]), this literature does not focus on the implications of the social network topology and interaction structure on information dissemination and belief formation. Most of the work relies on the assumption of perfect observability of the ordered history. Under this (implausible) assumption, the posterior beliefs form a martingale, which significantly simplifies the analysis, as one can simply apply Doob's martingale convergence theorem.

excessively influential group of individuals, then the social network has *nonexpanding observations*, and conversely, if there exists no excessively influential group, the network has *expanding observations*. This definition implies that most reasonable social networks have expanding observations, and in particular, a minimum amount of “arrival of new information” in the social network is sufficient for the expanding observations property. For example, the environment studied in most of the previous work in this area, where all past actions are observed, has expanding observations. Similarly, a social network in which each individual observes one uniformly drawn individual from those who have taken decisions in the past or a network in which each individual observes his immediate neighbor all feature expanding observations. Note also that a social network with expanding observations need not be *connected*. A simple, but typical, example of a network with nonexpanding observations is the one in which all future individuals only observe the actions of the first $K < \infty$ agents.

We establish the following result for the perfect Bayesian Nash Equilibrium of the learning game:

Theorem 3 (Impact of network topology on learning) *If the network topology is nonexpanding, then there will not be asymptotic learning. Conversely, if private beliefs are unbounded and the network topology is expanding, then there will be asymptotic learning.*

This is a striking result (particularly if we consider unbounded beliefs to be a better approximation to reality than bounded beliefs), since, as explained above, almost all reasonable social networks have the expanding observations property. This theorem, for example, implies that when some individuals, such as “informational leaders,” are overrepresented in the neighborhoods of future agents (and are thus “influential,” though not excessively so), learning may slow down, but asymptotic learning will still obtain as long as private beliefs are unbounded.

2.3 Value of Information in Shortest Path and Network Flow Optimizations

We now address the performance of decision networks by considering the limitations of *one agent* who makes a single decision under uncertainty [?]. The framework corresponds to, for example, a central decision maker who obtains information from either one or many imperfect distributed sensors, and, hence, has immediate applications in strategic planning under uncertainty. A natural question that arises in this setting is how information quantity (as determined by the quality or number and/or sensors) impacts decision quality.

In this framework of a single decision, we measure the performance of the agent’s decision, given by its state $x(1)$ at time $t = 1$ ($x(0)$ will be irrelevant). X is the set of decisions available to the agent so that $x(1) \in X$, and $l(x, W)$ is the cost (performance) of a decision $x \in X$ where the additional argument W (the state of the world) acts as a random perturbation of the decision’s quality. The agent’s private information Y about W dictates the action. The measure of performance for information governed by the distribution p_{WY} is

$$J(p_{WY}) = E \left[\min_{x(1) \in X} E[l(x(1), W)|Y] \right].$$

We continue the general setting stated previously in a performance-centric problem of high importance: shortest path optimization on a graph. In this problem, the decision set X is the set of possible paths in a directed acyclic graph $\mathcal{G} = (V, E)$ with vertices V and edges E , and each edge $e \in E$ has a random edge weight W_e .

Despite the computation of $J(p_{WY})$ being NP-Hard, we can leverage the geometric properties of the shortest-path polytope X of \mathcal{G} and the properties of $\mathcal{P}(c)$ to obtain fundamental performance bounds as well as simple characterizations of the actionable information [?]. First, rather than applying a traditional Shannon channel capacity parameterization of imperfect information, we adopt a different representation. Specifically, we use sets $\mathcal{P}(c)$ parameterized by a scalar c , which we still term *capacity*, according to $\mathcal{P}(c) = \{p_{WY} \mid \text{VAR}[E[W|Y]] \leq c\}$. The optimizing distribution $p_{WY}^* \in \mathcal{P}(c)$ as well as the best achievable performance from a c -amount of information is determined by the solution to

$$J(c) = \min_{p_{WY} \in \mathcal{P}(c)} J(p_{WY}). \quad (3)$$

Using this definition for information allows us to quantify the value of information according to the following theorem [?].

Theorem 4 (Value of information) *A lower bound for shortest path optimization performance under capacity c is*

$$J(c) \geq J(0) - \frac{d}{2}\sqrt{c},$$

where d is the diameter of X . The bound is “sharp” if $E[W] = 0$ and $C \leq \text{VAR}[W_e]$ for all $e \in E$.

In short, the fastest rate of improvement for shortest path optimization is the square root of the capacity c ³. We can further leverage the geometry of X to characterize the set actionable information that should be communicated to the agent.

Theorem 5 (Actionable information) *The actionable component of the information vector \hat{W} is that which lies in the smallest subspace containing X .*

The theorem tell us information contained in this subspace is all that is needed to choose the optimal path. The orthogonal component only improves estimation power, which is irrelevant to the agent’s objective. In fact, as long as the variance of the estimate in the actionable subspace is c , the estimate itself can be arbitrarily bad, and the agent can still achieve the same performance.

In [?], it is shown that a practical scheme for concentrating the information vector to its actionable component is to simply compare two paths of the graph. In the Gaussian case, this choice provably provides good performance.

³We can extend the same bound to any linear, combinatorial problem like shortest path optimization, and, further, so long as the original combinatorial problem can be solved in polynomial time, the lower bound can be computed in polynomial time (since d must be computed). The bound also holds for affine costs of the form $l(x, W) = g(x) + x^T W$.

Theorem 6 *If we restrict $p_{WY} \in \mathcal{P}'(c)$ where*

$$\mathcal{P}'(c) = \mathcal{P}(c) \cap \{p_{WY} | W \text{ and } Y \text{ are jointly Gaussian}\},$$

then

$$J(0) - \frac{d}{2} \sqrt{\frac{2}{\pi}} \sqrt{c} \leq J(c) \leq J(0) - \sqrt{\frac{1}{2\pi}} \sqrt{c}$$

where the optimizing $p_{WY}^ \in \mathcal{P}'(c)$ places all sensors along two paths of the graph.*

This performance-centric, information-energy setting can further be extended to a quasi-dynamic setting where information is gradually revealed to the agents. If information is randomly broadcast to the agent over time with no consideration to the agent's past decision, one can show that $J(c) \geq J(0) - \frac{d}{2} \sqrt{c - \Delta}$ where Δ is related to the co-variance of the information and how the dimensionality of X is reduced as decisions are made (which impacts how information energy must be concentrated). It can also be shown that the fundamental limit can be improved if future information accounts for the agent's past decisions.

We can immediately generalize the single-agent shortest path problem to multi-agent network flow optimization. In this problem, we have R agents who seek to traverse a graph along paths of minimal length. However, the length of each edge is determined by both a random length W as well as congestion due to multiple agents trying to access the edge simultaneously.

Formally, we extend the decision set to X^R , where X is the set of paths in the graph, and X^R is the set of paths that each of the possible R agents can take. A natural definition for the performance of the distributed system is the cumulative length taken by all agents through the graph, denoted $l(X^R, W)$.

As in the shortest path case, determining the value of information is computationally hard, but we can derive a fundamental bound for the value of information by leveraging geometric properties of \tilde{X} . This bound is given by the following theorem [?].

Theorem 7 (Value of information in network flow optimization) *A lower bound for network flow optimization among cooperative, distributed agents under capacity c is $J(c) \geq J(0) - O(c)$.*

Simulations verify that performance can improve linearly in the case of i.i.d. Gaussian edge weights and $c \leq \text{VAR}[W_e]$. Now, unlike shortest path optimization, the set of actionable information is no longer the subspace containing X because the relative number of agents to the variances of the individual edges impacts performance⁴. However, it is possible to show that performance does improve at the rate \sqrt{c} if information is concentrated to at most two paths, consistent with shortest path optimization.

3 Dynamic Decision Networks under Communication Constraints

We now consider a dynamic setting where stability and robustness must be considered in addition to performance. We will see that graph topology, information rate, and the rate at which agents respond to information can all significantly influence these important factors.

⁴This is due, in part, to the fact that the $l(X^R, W)$ is a quadratic function, not a linear function as in shortest path optimization

3.1 Stability of Capacity-Constrained Computation Network

Let us assume that an efficient distributed iterative algorithm exists for the computation of a certain function on a network whose links are assumed to support the noiseless transmission of real-valued messages. Important examples include iterative distributed averaging algorithms, as well as belief propagation algorithms for the computation of marginal distributions in graphical models. When dealing with the finite-capacity, and possibly noisy, communication link setting outlined in the previous section, a natural approach consists in trying to adapt such algorithms in order to cope with imperfect information transmission on the channels. Provided that the original iterative algorithm can be seen as a contraction in some metric space, it is not hard to see that the addition of bounded noise in each iteration would result in the accumulation of some bounded noise, whose magnitude can be controlled by controlling the magnitude of the noise introduced in each iteration.

However, typically distributed algorithms are not contractions globally, but rather exhibit some marginally stable, or unstable subspaces. As an example, typical iterative distributed averaging algorithms can be represented as multiplication by some irreducible doubly stochastic matrix P . Such matrices have a one-dimensional marginally stable eigenspace generated by the all-1 vector, and are contractions when restricted as operators to the orthogonal subspace to such eigenspace. When the noisy channel transmission allows for perfect feedback, whereby every node has knowledge of the corrupted message that its neighbors receive from itself (a particular case of this is quantized transmission), then it can compensate for the noise by subtracting it from its current state. Formally, if $x_i(t)$ is agent i 's state at time t , and $m_{ij}(t)$ is j 's estimate of $x_i(t)$, then each node simultaneously updates its state as

$$x_i(t+1) = \sum_j P_{ij} m_{ij}(t) - \sum_j P_{ji} (m_{ji}(t) - x_i(t)).$$

This update rule is such that the average of the states $n^{-1} \sum_i x_i(t)$ is preserved, so that noise acts only on the subspace orthogonal to the all-1 vector, where the algorithms is contractive. Indeed, one has the following result [?].

Theorem 8 (Effect of information concentration on stability) *Let ρ be the essential spectral radius of P , i.e. the second largest modulus of eigenvalues. Assume that $E[(x_i(t) - m_{ij}(t))^2] \leq \alpha^2$ for all i, j and $t \geq 0$. Let $\bar{x} = n^{-1} \sum_j x_j(0)$. Then*

$$\limsup_{t \rightarrow +\infty} \sum_i E[(x_i(t) - \bar{x})^2] \leq \frac{\alpha^2}{1 - \rho^2}$$

3.2 Stability and Performance of Capacity-Constrained Feedback Control

One of the most powerful results capturing performance trade-offs in a stable feedback system is the Bode integral formula [?]. In this classical result, it can be shown that for any strictly proper LTI plant P with unstable poles $\{\lambda_i\}_i$, the transfer function $S(z) = \frac{e(z)}{d(z)}$ between the disturbance d and the input e to P (also known as the *sensitivity function*), must satisfy the constraint:

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} [\log |S(e^{j\omega})|]_- d\omega + \frac{1}{2\pi} \int_{-\pi}^{\pi} [\log |S(e^{j\omega})|]_+ d\omega = \sum \max\{0, \log(|\lambda_i|)\} \quad (4)$$

where $[\log |S(e^{j\omega})|]_- = \min\{0, \log |S(e^{j\omega})|\}$ and $[\log |S(e^{j\omega})|]_+ = \max\{0, \log |S(e^{j\omega})|\}$.

The constraint implies that the sensitivity cannot be small at all frequencies, i.e., a reduction in $\int_{-\pi}^{\pi} [\log |S(e^{j\omega})|]_- d\omega$ is achieved at the expense of increasing $\int_{-\pi}^{\pi} [\log |S(e^{j\omega})|]_+ d\omega$.

A natural question in this setting is: Under what conditions can we break the Bode Integral Formula or must it always hold? We addressed this question in [?] by analyzing the information dynamics of a feedback system using an entropy-flow analysis. First, we take an information-theoretic view of the disturbance d and controller signal e as corrupted messages from K and P and analyze how restrictions on $\{K, P\} \in \mathcal{F}$ defined by a set \mathcal{F} of allowable systems impacts the information content of these messages.

Specifically, let $d(t) = W(t)$ be a dynamic state-of-the-world acting as a disturbance on P , and let $e = d - m_{KP}$. Further define

$$\begin{aligned}\mathcal{M} &= \{\{m_{PK}, m_{KP}\} | m_{PK} \text{ is a linear function of } P\text{'s internal state}\} \\ \mathcal{F}_c &= \{\{K, P\} | K \text{ is causal and } P \text{ is linear, controllable}\}.\end{aligned}$$

We represent causality in the entropy-flow analysis by the flow constraint

$$I(d(t); (P_t u, P_0 x) | P_{t-1} d) = 0 \quad \forall t \geq 0$$

where P_t is the projection operator defined as $P_t a = (a(0), \dots, a(t), 0, \dots)$ for any signal $a = (a(t))_t$. Essentially, the constraint implies that the controller provides no information about future signals given the past.

The assumption of causality immediately yields an interesting performance limitation: the entropy of the input e into P cannot be decreased below the external entropy injected into the system, formalized by the following theorem.

Theorem 9 *For any $\{K, P\} \in \mathcal{F}_c$ and $\{m_{PK}, m_{KP}\} \in \mathcal{M}$, $h(e(t)) \geq h(d(t)) + I(x(0); P_t e)$ for all $t \geq 0$.*

This limitation is independent of stability and the function of K (linear, non-linear, finite alphabet, etc...). It is also the basis for other performance limitations yielded through applying additional assumptions K . In particular, define

$$\mathcal{F}_{cs} = \{\{K, P\} | K \text{ is causal, stabilizing and } P \text{ is linear, controllable}\}.$$

We can represent stability in the entropy flow framework as a constraint on the variance of the state:

$$\sup_t E [x^T(t)x(t)] \leq \infty.$$

Also assume that d and e are asymptotically stationary stochastic signals (the weakest assumption under which they have power spectral densities). Let \hat{F}_d and \hat{F}_e be the respective power spectral densities of d and e , and define $S(\omega) = \sqrt{\hat{F}_e(\omega)/\hat{F}_d(\omega)}$, a direct generalization of the sensitivity function to a stochastic setting. Under these mild assumptions, we get the following theorem.

Theorem 10 (Fundamental limit of causal feedback) *For any $\{K, P\} \in \mathcal{F}_{cs}$ and $\{m_{PK}, m_{KP}\} \in \mathcal{M}$, the Bode Integral Formula holds.*




Figure 1: Shannon-Bode Tradeoffs: White Area Depends on Information Rate

The most interesting aspect of this theorem is that the form of K (linear, non-linear, lookup table, finite alphabet output, etc...) is irrelevant, and thus it yields a fundamental limit for any feedback system. The information theoretic analysis used to re-derive these limitations allowed us to derive performance limitations that cannot be addressed using traditional analysis techniques [?].

Theorem 11 (Generalized Bode Integral) *If $\{K, P\} \in \mathcal{F}_{cs}$ and if \mathcal{M} is constrained so that m_{PK} passes through a noisy digital memoryless channel with capacity c , then*

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} [\log |S(e^{j\omega})|]_- d\omega \geq \sum \max\{0, \log(|\lambda_i|) - c\}. \quad (5)$$

The bound is tight for certain Gaussian channels. By means of an argument similar to the water-bed effect, the inequality (??) asserts a limitation on the maximum allowable disturbance rejection over any given bandwidth $\Delta\omega$ in terms of the unstable modes of P and the information rate c . If the information rate is unconstrained ($c = \infty$), there is no limitation on attenuation other than, of course, the classical Bode Integral Formula.

Another information-theoretic analysis yielded a new generalization of the Bode Integral Formula to a case where it can be “broken” by giving the controller limited access to d by means of an early warning system [?]. In this context, d is now assumed to be a filtered white noise process w passing through a shaping filter G , and it is assumed that it takes $m > 0$ time units to reach P . The early warning system has access to d without the delay and uses a channel (wireless, for example) to send d to K . In our general framework, this is equivalent to K having some observation $Y_K(t)$ of $d(t)$. The new fundamental performance limitation we derived described in the following theorem.

Theorem 12 (Value of lookahead information) *If $\{K, P\} \in \mathcal{F}_{cs}$, $\{m_{PK}, m_{KP}\} \in \mathcal{M}$, and if K has access to an early-warning system with capacity c (that is, $I(Y_K(t); d(t)) \leq c$), then*

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} \log |S(e^{j\omega})| d\omega \geq \sum \max\{0, \log(|\lambda_i|)\} - c \quad (6)$$

In the case of an additive white Gaussian d and $m > 1$, the bound is tight.

A direct interpretation of this limitation is that disturbance rejection improved *linearly* with information rate.

3.3 Stability and Performance in Congestion Games with Learning

As transportation demand is fast approaching its infrastructure capacity, social planning for efficient usage of transportation networks (TNs) is attracting renewed research interest. Recent technological advancements are making available intelligent traveler information systems which have the capability to provide real-time, location specific traffic information and recommendations to the drivers and thereby enabling them to re-plan their routes during their trip. Transportation networks provide a prototypical example of distributed decision networks. Moreover, their analysis and control is made particularly challenging by the fact that they involve huge numbers of self-interested, bounded-rational agents whose behavior may be controlled only indirectly through incentives. In fact, transportation networks have been investigated in the economics literature, most notably in the context of (learning and evolution in) congestion games. However, these approaches tend to neglect most of the physical aspects of traffic dynamics, and are therefore unable to explain, e.g., transient behaviors occurring in response to sudden, and possibly disruptive, changes in the network characteristics.

In the works [?, ?], we proposed a novel framework for the analysis of stability and robustness properties of traffic networks. In our model, we abstract the topology of the transportation network by a directed acyclic graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, in which each directed edge $e = (v, v')$ represents a road, and each node $v \in \mathcal{V}$ represents a junction. We assume that there is a single origin-destination pair, $v_o, v_d \in \mathcal{E}$, with a constant unitary incoming flow in v_o , and that the size of the drivers population is so large to be efficiently approximated by a continuum of agents. Then, we consider a system whose state consists in: a probability vector $\pi(t) = \{\pi_p(t) : p \in \mathcal{P}\}$ over the set \mathcal{P} of simple paths from origin to destination, which takes into account the fraction of agents preferring a path with respect to the others; and a vector $\rho(t) = \{\rho_e(t) : e \in \mathcal{E}\}$ whose components correspond to the car density on the different roads. The car density vector $\rho(t)$ evolves as the drivers, modeled as boundedly rational agents, navigate their way through the network by combining their preference toward the different paths with the observation of the current local congestion levels in the network. On the other hand, the vector $\pi(t)$ evolves as the agents adapt their preferences toward the different paths using global information on the current congestion levels on the network. We assume that such global information is available at a time scale much slower than the typical time scale at which the actual drivers' dynamics occur. In this way, the dynamics of $\pi(t)$ and $\rho(t)$ becomes intertwined through two feedback loops, each involving a significantly different kind of information: local information at a fast scale, and global information at a slow scale. The system dynamics is then described by a system of ordinary differential equations

$$\begin{aligned} \frac{d}{dt}\rho_e &= f_v^- G_e(\rho^v, \pi) - f_e, & e \in \mathcal{E}, \\ \frac{d}{dt}\pi_p &= \eta (F_p(\rho) - \pi_p), & p \in \mathcal{P}, \end{aligned}$$

where $f_e = f_e(\rho_e)$ is the flow on link e , modeled as an increasing function of the traffic density with a maximum flow capacity C_e ; $\lambda_v := 1$ if $v = v_o$, $\lambda_v^- = \sum_{e \in \mathcal{E}_v^-} f_e$ if $v \neq v_o$, is the incoming flow in node v ; $\rho^v = \{\rho_e : e \in \mathcal{E}_v^+\}$ is the vector of traffic densities on the out-going edges of node v ; $G_e(\rho^v, \pi)$ is the fraction of drivers taking edge e when crossing node v , when the local density is ρ^v , and the path preferences profile is π ; $\eta > 0$ is the ratio (typically small) between the characteristic times of the fast and slow dynamics; $F(\rho) = \min_{\tilde{\pi}} \left\{ \sum_p \sum_{e \in p} \tilde{\pi}_p t_e(\rho_e) + \beta H(\tilde{\pi}) \right\}$ is the noisy best response to the current traffic density on the network, with $\beta > 0$, $H(\tilde{\pi})$ a convex function, and $t_e(\rho_e)$ the average delay on edge e .

A practical scenario to help envision this setup is where every driver is equipped with a smart navigation unit that recommends a direction to the driver based on its computations on global traffic information. Drivers augment this recommendation with the local information to navigate her way through the network. The resources to collect global information and compute optimal paths scale with the size of the network, and hence it is reasonable to expect that the navigation units will update their recommendations relatively infrequently as compared to typical transit times for large networks, and that the drivers are aware of this latency. Therefore the traffic dynamics are significantly influenced by the drivers response to local information.

In [?], we analyze the stability properties of this dynamical system. Under very mild assumptions on the drivers behavior, we show that system converges to a neighborhood of the Wardrop equilibrium. The latter is a well known notion of equilibrium configuration characterized by equal expected delay on every path from source to destination which is effectively chosen by some agent, whereby no agent has any incentive to switch. Formally, $\rho^* := \{\rho_e^* : e \in \mathcal{E}\}$ is a Wardrop equilibrium if $\sum_{e \in \mathcal{E}_v^-} f_e^* = \sum_{e \in \mathcal{E}_v^+} f_e^*$ for all $v \notin \{v_o, v_d\}$, and

$$\rho_e^* > 0, \quad \forall e \in p \quad \implies \quad \sum_{e \in p} t_e(\rho_e^*) \leq \sum_{e \in q} t_e(\rho_e^*), \quad \forall q \in \mathcal{P}.$$

Wardrop equilibria have been the object of big research effort, especially in relationship to their inefficiency in terms of social optimum, and the possibility to stir the Wardrop equilibrium towards a more socially efficient configuration through the use of tolls. In [?], it was shown that the asymptotic distance from the Wardrop equilibrium is controlled by both the time-scale ratio η and the noise level β :

Theorem 13 (Impact of information and control on convergence) *Let ρ^* be the unique Wardrop equilibrium of the network and $f^* = \{\mu_e(\rho_e^*) : e \in \mathcal{E}\}$ the corresponding flow vector. Then,*

$$\limsup_{t \rightarrow +\infty} \|f(t) - f^*\| \leq K(\eta + \beta),$$

for some positive constant K .

In [?], we study robustness properties of the system to sudden disruptions. Such disruptions are modeled as drastic changes in the physical properties of some of the links, which decrease (and possibly annihilate) their flow capacity. We then look at the evolution of the system assuming that it is started at the Wardrop equilibrium of the unperturbed system, and that the agents' global preference toward the different paths does not change significantly

(this is because $\pi(t)$ evolves at a time scale too slow for being significant in the presence of sudden disruptions). We characterize the margin of stability γ of the original equilibrium, i.e. the minimum total loss in flow capacity that makes the system unstable, as the minimum node cut of the network.

Theorem 14 (Robustness to graph perturbations) $\gamma = \min_{v \neq v_o} \sum_{e \in \mathcal{E}_v^+} C_e - f_e^* .$

This quantification of the margin of stability has to be contrasted to the min-cut capacity of the network. Indeed the former is equilibrium-dependent, always not larger than the latter, which is equilibrium independent, and in fact typically much smaller than it. Such a gap between the two is in fact a consequence of the locality of information available to the agents. The margin of stability provides a second order optimization parameter for the optimal choice of tolls.

References

- [1] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar. Bayesian learning in social networks. *Economic Reviews*, 2010.
- [2] A. Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107:797–817, 1992.
- [3] A. Banerjee and D. Fudenberg. Word-of-mouth learning. *Games and Economic Behavior*, 46:1–22, 2004.
- [4] S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as information cascades. *Journal of Political Economy*, 100:992–1026, 1992.
- [5] H. W. Bode. *Network Analysis and Feedback Amplifier Design*. D. Van Nostrand, Princeton, NJ, 1945.
- [6] S. Borade, B. Nakiboglu, and L. Zheng. Euclidean Information Geometry. *ISIT*, 2007.
- [7] G. Como and M. Dahleh. Lower bounds on the estimation error in problems of distributed computation. In *Proc. of ITA Workshop*, 2009.
- [8] G. Como, K. Savla, D. Acemoglu, M. Dahleh, and E. Frazzoli. Robust distributed routing in dynamical networks - part i: Locally responsive policies and weak resilience. In *IEEE Trans. AC*, 2013.
- [9] G. Como, K. Savla, D. Acemoglu, M. Dahleh, and E. Frazzoli. Robust distributed routing in dynamical networks - part ii: Strong resilience, equilibrium selection and cascaded failures. *IEEE Trans. AC*, 2013.
- [10] M.A. Dahleh G. Baldan. Channel optimization for binary hypothesis testing. In *LIDS Report*, 2013.
- [11] N.C. Martins and M.A. Dahleh. Feedback Control in The Presence of Noisy Channels: Bode-Like Fundamental Limitation of Performance. *IEEE Trans. on Automatic Control*.
- [12] N.C. Martins, M.A. Dahleh, and J.C. Doyle. Fundamental Limitations of Disturbance Attenuation with Side Information. *IEEE Trans. on Automatic Control*, January 2007.
- [13] M. Rinehart and M.A. Dahleh. The Networked Control Systems. *A Chapter in the "Impact of Control Technology"*, 2011.
- [14] M. Rinehart and M.A. Dahleh. The Value of Side Information in Shortest Path Optimization. *IEEE Trans. Automatic Control*, 2011.
- [15] L. Smith and P. Sorensen. Rational social learning with random sampling. unpublished manuscript, 1998.

- [16] L. Smith and P. Sorensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.

1.

1. Report Type

Final Report

Primary Contact E-mail

Contact email if there is a problem with the report.

dahleh@mit.edu

Primary Contact Phone Number

Contact phone number if there is a problem with the report

617-253-3892

Organization / Institution name

Massachusetts Institute of Technology

Grant/Contract Title

The full title of the funded effort.

The Value of Information in Distributed Decision Networks

Grant/Contract Number

AFOSR assigned control number. It must begin with "FA9550" or "F49620" or "FA2386".

FA9550-11-1-0312

Principal Investigator Name

The full name of the principal investigator on the grant or contract.

Munther Dahleh

Program Manager

The AFOSR Program Manager currently assigned to the award

Frederick A. Leve, Ph.D.

Reporting Period Start Date

09/30/2014

Reporting Period End Date

09/29/2015

Abstract

Please see attached report.

Distribution Statement

This is block 12 on the SF298 form.

Distribution A - Approved for Public Release

Explanation for Distribution Statement

If this is not approved for public release, please provide a short explanation. E.g., contains proprietary information.

SF298 Form

Please attach your [SF298](#) form. A blank SF298 can be found [here](#). Please do not password protect or secure the PDF. The maximum file size for an SF298 is 50MB.

[sf298_FA9550-11-1-0312.pdf](#)

Upload the Report Document. File must be a PDF. Please do not password protect or secure the PDF. The maximum file size for the Report Document is 50MB.

[Dahleh_final report_FA9550-11-1-0312pdf.pdf](#)

Upload a Report Document, if any. The maximum file size for the Report Document is 50MB.

Archival Publications (published) during reporting period:

Changes in research objectives (if any):

Change in AFOSR Program Manager, if any:

Extensions granted or milestones slipped, if any:

AFOSR LRIR Number

LRIR Title

Reporting Period

Laboratory Task Manager

Program Officer

Research Objectives

Technical Summary

Funding Summary by Cost Category (by FY, \$K)

	Starting FY	FY+1	FY+2
Salary			
Equipment/Facilities			
Supplies			
Total			

Report Document

Report Document - Text Analysis

Report Document - Text Analysis

Appendix Documents

2. Thank You

E-mail user

Mar 02, 2016 10:31:45 Success: Email Sent to: dahleh@mit.edu